

[Subscribe](#) (Full Service) [Register](#) (Limited Service, Free) [Login](#)Search: ☒ The ACM Digital Library ☐ The Guide

HTTP and name-value pairs and abstraction layer



THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)Terms used [HTTP](#) and [name value pairs](#) and [abstraction layer](#)

Found 48,627 of 132,85

Sort results
by

relevance

Display
results

expanded form

[Save results to a Binder](#)[Search Tips](#)☐ Open results in a new
windowTry an [Advanced Search](#)Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐

1 [Stateful distributed interposition](#)

John Reumann, Kang G. Shin

February 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 1Full text available: [pdf\(833.84 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Interposition-based system enhancements for multitiered servers are difficult to build because important system context is typically lost at application and machine boundaries. For example, resource quotas and user identities do not propagate easily between cooperating services that execute on different hosts or that communicate with each other via intermediary services. Application-transparent system enhancement is difficult to achieve when such context information is obscured by complex service ...

Keywords: Distributed computing, component services, distributed context, multitiered services, operating systems, server consolidation

2 [Extending Java for high-level Web service construction](#)

Aske Simon Christensen, Anders Møller, Michael I. Schwartzbach

November 2003 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 25 Issue 6Full text available: [pdf\(947.02 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We incorporate innovations from the <bigwig> project into the Java language to provide high-level features for Web service programming. The resulting language, JWIG, contains an advanced session model and a flexible mechanism for dynamic construction of XML documents, in particular XHTML. To support program development we provide a suite of program analyses that at compile time verify for a given program that no runtime errors can occur while building documents or receiving form input, and ...

Keywords: Interactive Web services, XML, data-flow analysis

3 Principled design of the modern Web architecture

Roy T. Fielding, Richard N. Taylor

May 2002 **ACM Transactions on Internet Technology (TOIT)**, Volume 2 Issue 2

Full text available:  pdf(335.47 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The World Wide Web has succeeded in large part because its software architecture has been designed to meet the needs of an Internet-scale distributed hypermedia application. The modern Web architecture emphasizes scalability of component interactions, generality of interfaces, independent deployment of components, and intermediary components to reduce interaction latency, enforce security, and encapsulate legacy systems. In this article we introduce the Representational State Transfer (REST) arc ...

Keywords: Network-based applications, REST, World Wide Web

4 Generating presentation constraints from rhetorical structure

Lloyd Rutledge, Brian Bailey, Jacco van Ossenbruggen, Lynda Hardman, Joost Geurts

May 2000 **Proceedings of the eleventh ACM on Hypertext and hypermedia**

Full text available:  pdf(233.19 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: authoring, constraints, meta-structure, presentation generation, rhetorics

5 The webspace method: on the integration of database technology with multimedia retrieval

Roelof van Zwol, Peter M. G. Apers

November 2000 **Proceedings of the ninth international conference on Information and knowledge management**

Full text available:  pdf(208.62 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: concept-based search, content-based information retrieval, daemon data dictionary, modelling data on the web

6 Abstraction-based intrusion detection in distributed environments

November 2001 **ACM Transactions on Information and System Security (TISSEC)**, Volume 4 Issue 4

Full text available:  pdf(590.61 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

Abstraction is an important issue in intrusion detection, since it not only hides the difference between heterogeneous systems, but also allows generic intrusion-detection models. However, abstraction is an error-prone process and is not well supported in current intrusion-detection systems (IDSs). This article presents a hierarchical model to support attack specification and event abstraction in distributed intrusion detection. The model involves three concepts: *system view*, *signature* ...

Keywords: Cooperative information systems, heterogeneous systems, intrusion detection, misuse detection

7 HydroJ: object-oriented pattern matching for evolvable distributed systems

Keunwoo Lee, Anthony LaMarca, Craig Chambers

October 2003 **ACM SIGPLAN Notices , Proceedings of the 18th ACM SIGPLAN conference on Object-oriented programing, systems, languages, and applications**, Volume 38 Issue 11

Full text available:  pdf(311.06 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


In an evolving software system, components must be able to change independently while remaining compatible with their peers. One obstacle to independent evolution is the *brittle parameter problem*: the ability of two components to communicate can depend on a number of *inessential* details of the types, structure, and/or contents of the values communicated. If these details change, then the components can no longer communicate, even if the *essential* parts of the message remain ...

Keywords: HydroJ, XML, distributed systems, dynamic dispatch, object-oriented programming, pattern matching, semi-structured data, software evolution, ubiquitous computing

8 A comparative study of language support for generic programming

Ronald Garcia, Jaakko Jarvi, Andrew Lumsdaine, Jeremy Siek, Jeremiah Willcock

October 2003 **ACM SIGPLAN Notices , Proceedings of the 18th ACM SIGPLAN conference on Object-oriented programing, systems, languages, and applications**, Volume 38 Issue 11

Full text available:  pdf(237.38 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Many modern programming languages support basic generic programming, sufficient to implement type-safe polymorphic containers. Some languages have moved beyond this basic support to a broader, more powerful interpretation of generic programming, and their extensions have proven valuable in practice. This paper reports on a comprehensive comparison of generics in six programming languages: C++, Standard ML, Haskell, Eiffel, Java (with its proposed generics extension), and Generic C. By implementi ...

Keywords: C#, C++, Eiffel, Haskell, Java, generic programming, generics, polymorphism, standard ML

9 Business-to-business interactions: issues and enabling technologies

B. Medjahed, B. Benatallah, A. Bouguettaya, A. H. H. Ngu, A. K. Elmagarmid

May 2003 **The VLDB Journal — The International Journal on Very Large Data Bases,**

Volume 12 Issue 1

Full text available:  [pdf\(558.34 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Business-to-Business (B2B) technologies pre-date the Web. They have existed for at least as long as the Internet. B2B applications were among the first to take advantage of advances in computer networking. The Electronic Data Interchange (EDI) business standard is an illustration of such an early adoption of the advances in computer networking. The ubiquity and the affordability of the Web has made it possible for the masses of businesses to automate their B2B interactions. However, several issues ...

Keywords: B2B Interactions, Components, E-commerce, EDI, Web services, Workflows, XML

10 Languages & Authoring for the Semantic Web: The Yin/Yang web: XML syntax and RDF semantics

Peter Patel-Schneider, Jérôme Siméon


May 2002 **Proceedings of the eleventh international conference on World Wide Web**Full text available:  [pdf\(162.67 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

XML is the W3C standard document format for *writing* and *exchanging* information on the Web. RDF is the W3C standard model for *describing the semantics* and *reasoning about* information on the Web. Unfortunately, RDF and XML---although very close to each other---are based on two different paradigms. We argue that in order to lead the Semantic Web to its full potential, the syntax and the semantics of information needs to work together. To this end, we develop a model-theory ...

Keywords: RDF, XML, data models, model theory, semantic web

11 HTTP Cookies: Standards, privacy, and politics

David M. Kristol

November 2001 **ACM Transactions on Internet Technology (TOIT)**, Volume 1 Issue 2Full text available:  [pdf\(390.38 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

How did we get from a world where cookies were something you ate and where "nontechnies" were unaware of "Netscape cookies" to a world where cookies are a hot-button privacy issue for many computer users? This article describes how HTTP "cookies" work and how Netscape's original specification evolved into an IETF Proposed Standard. I also offer a personal perspective on how what began as a straightforward technical specification turned into a political flashpoint when it tried to address nontechnical ...

Keywords: Cookies, HTTP, World Wide Web, privacy, state management

12 An architecture for secure wide-area service discovery

Todd D. Hodes, Steven E. Czerwinski, Ben Y. Zhao, Anthony D. Joseph, Randy H. Katz
March 2002 **Wireless Networks**, Volume 8 Issue 2/3

Full text available:  [pdf\(365.68 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The widespread deployment of inexpensive communications technology, computational resources in the networking infrastructure, and network-enabled end devices poses an interesting problem for end users: how to locate a particular network service or device out of hundreds of thousands of accessible services and devices. This paper presents the architecture and implementation of a secure wide-area Service Discovery Service (SDS). Service providers use the SDS to advertise descriptions of available ...

Keywords: location services, name lookup, network protocols, service discovery

13 Workshop and conference summaries: Exchange format bibliography

Holger M. Kienle
January 2001 **ACM SIGSOFT Software Engineering Notes**, Volume 26 Issue 1

Full text available:  [pdf\(616.88 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

This paper gives a brief bibliographical overview of exchange formats and related research areas. We classify exchange formats and try to give a brief assessment of the more interesting ones.

Keywords: Exchange format, bibliography, graph format, overview

14 Lightweight object-oriented shared variables for distributed applications on the Internet

Jacob Harris, Vivek Sarkar
October 1998 **ACM SIGPLAN Notices , Proceedings of the 13th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications**, Volume 33 Issue 10

Full text available:  [pdf\(1.78 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


This paper describes a lightweight yet powerful approach for writing distributed applications using shared variables. Our approach, called SHAREHOLDER, is inspired by the flexible and intuitive model of information access common to the World Wide Web. The distributed applications targeted by our approach all share a weak consistency model and loose transaction semantics, similar to a user's model of accessing email, bulletin boards, chat rooms, etc. on the Internet. The SHA ...

Keywords: distributed and parallel systems, language design and implementation, persistence, world wide web

15 Research sessions: distributed systems: Proxy-based acceleration of dynamically generated content on the world wide web: an approach and implementation

Anindya Datta, Kaushik Dutta, Helen Thomas, Debra VanderMeer, Suresha, Krithi Ramamritham

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available:  pdf(1.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

As Internet traffic continues to grow and web sites become increasingly complex, performance and scalability are major issues for web sites. Web sites are increasingly relying on dynamic content generation applications to provide web site visitors with dynamic, interactive, and personalized experiences. However, dynamic content generation comes at a cost --- each request requires computation as well as communication across multiple components. To address these issues, various dynamic content caching ...

Keywords: dynamic content, edge caching, proxy-based caching

16 Principled design of the modern Web architecture

Roy T. Fielding, Richard N. Taylor

June 2000 **Proceedings of the 22nd international conference on Software engineering**

Full text available:  pdf(217.34 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The World Wide Web has succeeded in large part because its software architecture has been designed to meet the needs of an Internet-scale distributed hypermedia system. The modern Web architecture emphasizes scalability of component interactions, generality of interfaces, independent deployment of components, and intermediary components to reduce interaction latency, enforce security, and encapsulate legacy systems. In this paper, we introduce the Representational State Transfer ...

Keywords: WWW, software architectural style, software architecture

17 Web technologies and applications (WTA): Migration to web services oriented architecture: a case study

Jia Zhang, Jen-Yao Chung, Carl K. Chang

March 2004 **Proceedings of the 2004 ACM symposium on Applied computing**

Full text available:  pdf(159.22 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The rapid emerging of web-services technology is dramatically changing the scenario of web application design and development. This paper presents a web-services oriented architecture. As a case study, the paper reports on an on-going project on the design and development of a pass-through authentication web-services for on-line electronic payment applications. This is a first step towards an electronic payment web-service.

Keywords: Web application development, case study, software architecture, web services

oriented architecture

18 The many faces of publish/subscribe

Patrick Th. Eugster, Pascal A. Felber, Rachid Guerraoui, Anne-Marie Kermarrec

June 2003 **ACM Computing Surveys (CSUR)**, Volume 35 Issue 2

Full text available:  pdf(456.46 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Well adapted to the loosely coupled nature of distributed interaction in large-scale applications, the publish/subscribe communication paradigm has recently received increasing attention. With systems based on the publish/subscribe interaction scheme, subscribers register their interest in an event, or a pattern of events, and are subsequently asynchronously notified of events generated by publishers. Many variants of the paradigm have recently been proposed, each variant being specifically adap ...

Keywords: Distribution, interaction, publish/subscribe

19 Information agents for automated browsing

Chanda Dharap, Martin Freeman

November 1996 **Proceedings of the fifth international conference on Information and knowledge management**

Full text available:  pdf(1.12 MB) Additional Information: [full citation](#), [references](#), [index terms](#)

20 Interoperable Web services for computational portals

Marlon Pierce, Geoffrey Fox, Choonhan Youn, Steve Mock, Kurt Mueller, Ozgur Balsoy

November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**

Full text available:  pdf(278.00 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Computational web portals are designed to simplify access to diverse sets of high performance computing resources, typically through an interface to computational Grid tools. An important shortcoming of these portals is their lack of interoperable and reusable services. This paper presents an overview of research efforts undertaken by our group to build interoperating portal services around a Web Services model. We present a comprehensive view of an interoperable portal architecture, beginning w ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)